



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-PROC-222915

Sequential Monte-Carlo Based Framework for Dynamic Data-Driven Event Reconstruction for Atmospheric Release

G. Johannesson, K.M. Dyer, W.G. Hanley, B. Kosovic,
S.C. Larsen, G.A. Loosmore, J.K. Lundquist, A.A. Mirin

July 17, 2006

Nonlinear Statistical Signal Processing Workshop
Cambridge, United Kingdom
September 13, 2006 through September 15, 2006

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

SEQUENTIAL MONTE-CARLO FRAMEWORK FOR DYNAMIC DATA-DRIVEN EVENT RECONSTRUCTION FOR ATMOSPHERIC RELEASE

*Gardar Johannesson, Kathleen M. Dyer, William G. Hanley, Branko Kosovic,
Shawn C. Larsen, Gwendolen A. Loosmore, Julie K. Lundquist, Arthur A. Mirin*

Lawrence Livermore National Laboratory, Livermore, CA, USA

ABSTRACT

The release of hazardous materials into the atmosphere can have a tremendous impact on dense populations. We propose an atmospheric event reconstruction framework that couples observed data and predictive computer-intensive dispersion models via Bayesian methodology. Due to the complexity of the model framework, a sampling-based approach is taken for posterior inference that combines Markov chain Monte Carlo (MCMC) and sequential Monte Carlo (SMC) strategies.

1. INTRODUCTION

Atmospheric event reconstruction refers to the process of estimating the characteristics of an unknown release of a chemical or biological agent into the atmosphere (e.g., when? where? how much?) and predicting its current and future dispersion. The first part is accomplished by coupling together an atmospheric dispersion model and (relatively) sparse sensor-data to extract information about the unknown release input parameters of the dispersion model. Given a characterization of the release, the resulting dispersion can be predicted, providing valuable information for consequence management. A dynamic atmospheric event reconstruction refers to a constant revision of our state-of-knowledge of the unknown release, and its dispersion, as the event unfolds (in real or near-real time) and more data becomes available.

A team of scientists at Lawrence Livermore National Laboratory is implementing a general framework to carry out dynamic atmospheric event reconstruction. The approach is (hierarchical) Bayesian, coupling together observed data and a given dispersion model, along with prior knowledge about source parameters, model uncertainty, and data accuracy. The computational framework is implemented in a (Linux) cluster environment and consists of a posterior-sampler interacting with a dispersion model-server that can carry out multiple dispersion model runs in parallel.

2. PROBLEM SETUP

Let $t = 1, 2, \dots$ index the time-periods (time-points) for which we wish to update our model and let:

$\tau_t \equiv$ the end time-point of the t -th characterization.

$\mathcal{T}_t \equiv (\tau_{t-1}, \tau_t]$, the t -th time interval.

The time-points $\{\tau_t\}$ can either be fixed in advance or, what is more likely the case, dynamically chosen based on the availability and coverage of the incoming data.

The unknown atmospheric release can be due to one or more sources. Let

$\theta_{t,i} \equiv$ collection of parameters characterizing the i -th source in \mathcal{T}_t , and let $\theta_t \equiv \{\theta_{t,i}\}$.

For example, a simple characterization of a single explosive-type of a source is given by its location (\mathbf{x}), time of explosion (T), and the amount of material being released (m); $\theta = \{\mathbf{x}, T, m\}$. Other examples include a source with a fixed location but a time-dependent release rate (e.g., a “leak”) and a source with both time-dependent release location and rate (i.e., a moving source).

The impact of the release can be observed in various ways, for example by a network of sensor instruments. In that case, let:

$\mathbf{s}_j \equiv$ the location of the j -th sensor.

$c_{j,k} \equiv$ the k -th average concentration reported from the j -th sensor.

$\mathcal{T}_{j,k} \equiv (\tau_{j,k}^s, \tau_{j,k}^e) \equiv$ the time interval in which $c_{j,k}$ was measured over.

And we denote by

$\mathbf{d}_t \equiv \{c_{j,k} : (j,k) \in \mathcal{I}_t\}$, where,

$\mathcal{I}_t \equiv \{(j,k) : \tau_{t-1} < \tau_{j,k}^e \leq \tau_t\}$, indexes the new data available in t -th time period.

Further, let $\mathbf{d}_{1:t} \equiv \{\mathbf{d}_1, \dots, \mathbf{d}_t\}$.

An atmospheric dispersion computer model is used to relate the source parameters to the sensor data. The dispersion

model yields concentration predictions given various input parameters, including characterization of the emitting source. Let

$\hat{C}_{j,k} = \hat{C}_{j,k}(\boldsymbol{\theta}_{1:t}) \equiv$ the model-predicted average concentration at location \mathbf{s}_j in the time-period $\mathcal{T}_{i,k}$.

$\hat{\mathbf{C}}_t \equiv \hat{\mathbf{C}}_t(\boldsymbol{\theta}_{1:t^*}) \equiv \{\hat{C}_{j,k} : (j,k) \in \mathcal{I}_t\}$, $t \leq t^*$, the model-predicted concentrations corresponding to \mathbf{d}_t .

The chosen dispersion model is not perfect, but an approximation to the underlying physical dispersion processes. We therefore define:

$C_{j,k} = C(\mathbf{s}_j, \mathcal{T}_{j,k}) \equiv$ the true (unknown) average concentration at location \mathbf{s}_j over the time period $\mathcal{T}_{j,k}$.

$\mathbf{C}_t \equiv \{C_{j,k} : (j,k) \in \mathcal{I}_t\}$.

3. ON ATMOSPHERIC DISPERSION MODELS

A core component of an atmospheric event reconstruction is an efficient use of atmospheric dispersion models. These models range in complexity, from simple and fast (Gaussian) ‘‘puff’’ models, such as the INPUFF model [1] that we use later in Section 6, to more computationally demanding dispersion models, such as the Lagrangian-based LODI code [2]. For the more computationally demanding dispersion models, it is crucial to minimize the number of model-runs needed for source characterization. We now briefly describe two approaches that can help in this regard.

Some dispersion models yield predicted concentration levels that scale linear with the amount of material being released. If this is the case, then, for example for an explosive source $\boldsymbol{\theta} = \{\mathbf{x}, T, m\}$, it is sufficient to carry out a *single* model-run at a proposed explosion location and time (\mathbf{x}, T) to derive the predicted sensor concentrations $(\hat{C}_{j,k})$ for various release masses (m) .

To efficiently take advantage of the above linearity, the source parameters need to be ‘‘mapped’’ to a spatio-temporal grid before being used by a simulation-based dispersion model. For example, the proposed location and release time of an explosive source would be assigned to the closest source grid-point. This introduces the idea of spatio-temporal resolution. Thus, we extend the notation for the predicted concentration levels to include a resolution index R and let

$\hat{\mathbf{C}}_t^{(R)} \equiv$ the predicted sensor concentrations $\hat{\mathbf{C}}_t$ at a given source-resolution R .

4. MODEL DEVELOPMENT

Given sensor data from t time periods (batches), $\mathbf{d}_{1:t}$, our primary goal is to conduct inference on the source parameters, $\boldsymbol{\theta}_{1:t}$, via the posterior distribution

$$\pi_t(\boldsymbol{\theta}_{1:t}) \equiv p(\boldsymbol{\theta}_{1:t} | \mathbf{d}_{1:t});$$

the probability distribution of the source parameters conditional on the observed data. In addition, as our setup is dynamic, we seek a smooth transition from posterior inference at time $t - 1$ to time t ; from $\pi_{t-1}(\boldsymbol{\theta}_{1:t-1})$ to $\pi_t(\boldsymbol{\theta}_{1:t})$.

The source terms are not the only parameters of interest. We also seek posterior inference on the impact of the release, that is, on the resulting concentration levels,

$$\pi_t(C(\mathbf{u}, \tau)) \equiv p(C(\mathbf{u}, \tau) | \mathbf{d}_{1:t}), \quad 0 < \tau \leq \tau_t, \quad \mathbf{u} \in D,$$

where D is our spatial domain of interest. In general, we seek to have access to the joint posterior

$$\pi_t(\mathbf{C}_{1:t}, \boldsymbol{\theta}_{1:t}) \equiv p(\mathbf{C}_{1:t}, \boldsymbol{\theta}_{1:t} | \mathbf{d}_{1:t}).$$

When proceeding from time period $t - 1$ to t , the atmospheric event reconstruction problem has a natural hierarchical breakdown:

Data Model: A conditional probability distribution describing the variation in the newly available data, \mathbf{d}_t , given the true underlying concentrations and past data,

$$p(\mathbf{d}_t | \mathbf{C}_{1:t}, \mathbf{d}_{1:t-1}). \quad (1)$$

Process Model: A probability model describing the variation in the current concentration levels, \mathbf{C}_t , given model predictions and past concentration levels,

$$p(\mathbf{C}_t | \hat{\mathbf{C}}_{1:t}, \mathbf{C}_{1:t-1}) = p(\mathbf{C}_t | \boldsymbol{\theta}_{1:t}, \mathbf{C}_{1:t-1}), \quad (2)$$

where the second expression follows since $\hat{\mathbf{C}}_{t^*} = \hat{\mathbf{C}}_{t^*}(\boldsymbol{\theta}_{1:t})$, $t^* = 1, \dots, t$.

Parameter Model: A prior parameter model,

$$p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_{1:t-1}). \quad (3)$$

Before going further, we note that each of these models might have a collection of (hyper) parameters. For example, there might be some parameters that describe the size of the dispersion model-error (i.e., additional parameters associated with (2)). Jointly, we have

$$p(\mathbf{C}_t, \boldsymbol{\theta}_t | \mathbf{C}_{1:t-1}, \boldsymbol{\theta}_{1:t-1}) = p(\mathbf{C}_t | \mathbf{C}_{1:t-1}, \boldsymbol{\theta}_{1:t})p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_{1:t-1}).$$

The joint (prior) distribution of the model parameters can therefore be written as

$$p(\mathbf{C}_{1:t}, \boldsymbol{\theta}_{1:t}) = \prod_{t^*=1}^t p(\mathbf{C}_{t^*}, \boldsymbol{\theta}_{t^*} | \mathbf{C}_{1:t^*-1}, \boldsymbol{\theta}_{1:t^*-1}),$$

where we define $\mathbf{C}_{1:0} = \boldsymbol{\theta}_{1:0} = \emptyset$ (an empty set of parameters).

Due to the dynamic setup and the hierarchical breakdown within each time step, inference flows naturally from time period $t - 1$ to t . Assume at time $t - 1$ we have access to the joint posterior distribution of all parameters of interest,

$$\pi_{t-1}(\mathbf{C}_{1:t-1}, \boldsymbol{\theta}_{1:t-1}) = p(\mathbf{C}_{1:t-1}, \boldsymbol{\theta}_{1:t-1} | \mathbf{d}_{1:t-1}).$$

The dynamic parameter model (3) for the source terms yields the one-step-ahead predictive distribution as

$$\pi_{t-1}(\boldsymbol{\theta}_{1:t}) = p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_{1:t-1})\pi_{t-1}(\boldsymbol{\theta}_{1:t-1}).$$

More generally, jointly we have that,

$$\pi_{t-1}(\mathbf{C}_{1:t}, \boldsymbol{\theta}_{1:t}) = p(\mathbf{C}_t, \boldsymbol{\theta}_t | \mathbf{C}_{1:t-1}, \boldsymbol{\theta}_{1:t-1}) \times \pi_{t-1}(\mathbf{C}_{1:t-1}, \boldsymbol{\theta}_{1:t-1}).$$

Then, given (potentially) new data at time step t , the joint posterior at t is given by

$$\pi_t(\mathbf{C}_{1:t}, \boldsymbol{\theta}_{1:t}) \propto p(\mathbf{d}_t | \mathbf{C}_{1:t}, \mathbf{d}_{1:t-1})\pi_{t-1}(\mathbf{C}_{1:t}, \boldsymbol{\theta}_{1:t}). \quad (4)$$

5. POSTERIOR INFERENCE

For posterior inference we adopt a sequential Monte Carlo method. In designing a posterior sampler, we seek robustness, adaptiveness, and effective use of the (sometimes computational demanding) atmospheric dispersion prediction code. We now give a broad description of the sampler's design.

The sampler consists of three components: (1) Initialization; the generation of the initial sample, (2) Rejuvenation; the refinement and 'cooling' of the current sample, and (3) Augmentation; the extension of time-dependent parameters. The design is inspired by the use of annealing (bridging) methods (as in [3, 4]), the use of auxiliary variables (as in [5]), the use of MCMC kernels (as in [6, 4]), and by the adaptability of population Monte Carlo methods [7]. Many of these approaches are summarized, generalized, and extended by Del Moral et al. [8] under the name of "Sequential Monte Carlo Samplers". What is new, and not covered by above references, is the implementation of an adaptive annealing schedule and a constant refinement of the spatial resolution of the source parameters.

In what follow, we let $\mathbf{X}_t \equiv \{\boldsymbol{\theta}_t, \mathbf{C}_t\}$.

(1): Initialization

Given the initial data \mathbf{d}_1 at $t = 1$, we seek to draw an importance sample from the 'heated' and 'coarsened' posterior distribution

$$\pi_1^{(T_1, R_1)}(\mathbf{X}_1) \propto p^{(R_1)}(\mathbf{d}_1 | \mathbf{X}_1)^{1/T_1} p^{(R_1)}(\mathbf{X}_1), \quad (5)$$

where R_1 indexes the initial (coarse) spatial resolution of the source parameters (a provided input parameter) and $T_1 > 1$ is the initial annealing temperature. The algorithm used to generate the importance sample is as follows:

1. Generate $\mathbf{X}_1^{(i)} \sim p^{(R_1)}(\cdot)$; $i = 1, \dots, N$.
2. Compute $L_1^{(i)} \propto p^{(R_1)}(\mathbf{d}_1 | \mathbf{X}_1^{(i)})$.
3. Find a temperature $T_1 \geq 1$ that yields importance weights $w_1^{(i)} \propto (L_1^{(i)})^{1/T_1}$ with effective sample size (ESS) [9] just above a given threshold, say $N/2$.

Let $\{\mathbf{X}_1^{(i)}, w_1^{(i)}\}$ be the resulting importance sample.

(2): Rejuvenation

Let $\{\mathbf{X}_{1:t}^{(i)}, w_{1:t}^{(i)}\}$ be an importance sample from

$$\pi_t^{(T, R)}(\mathbf{X}_{1:t}) \propto p^{(R)}(\mathbf{d}_{1:t} | \mathbf{X}_{1:t})^{1/T} p^{(R)}(\mathbf{X}_{1:t}). \quad (6)$$

We use MCMC kernels to further 'cool' and refine the current sample using (broadly) the following steps:

1. Resample the current sample .
2. Select a new source-resolution index R^* using the empirical distribution of the current source parameters.
3. Adapt MCMC kernels using the current sample (e.g., the "step-sizes" in Metropolis-Hasting-type of a random-walk kernels) and carry out MCMC proposals.
4. Select a new temperature T^* , $1 \leq T^* < T$, that results in an expected MCMC acceptance ratio and ESS just above given thresholds.
5. Accept/reject M-H proposals using the new T^* in (6), yielding a MCMC-rejuvenated sample $\{\mathbf{X}_{1:t}^{*(i)}\}$
6. Compute the importance weights [8]

$$w_{1:t}^{*(i)} \propto \pi_t^{(T^*, R^*)}(\mathbf{X}_{1:t}^{*(i)}) / \pi_t^{(T, R)}(\mathbf{X}_{1:t}^{*(i)}).$$

This yields an importance sample $\{\mathbf{X}_{1:t}^{*(i)}, w_{1:t}^{*(i)}\}$ at temperature T^* and source-resolution R^* .

(3): Augmentation

Let $\{\mathbf{X}_{1:t}^{(i)}, w_{1:t}^{(i)}\}$ be the current importance sample from (6). Due to the inherent time delay in what happens at the source and what is observed at the sensors, any potential new data \mathbf{d}_{t+1} is mostly informative about the past state of the source. We therefore adopt auxiliary particle-filter techniques [5] in extending the current importance sample to time period $t + 1$. The only major deviation from the classical auxiliary approach is the adaptive selection of the annealing temperature, along the lines of the Initialization step above, except the temperature change is only applied to the new data.

Putting it All Together

In short, we carry out the Initialization step followed by multiple Rejuvenation steps (the number of Rejuvenations needed can be determined by monitoring changes in some population statistics from one iteration to the next; e.g., changes in mean and variance). The initial sample is then augmented using the Augmentation step, followed by multiple Rejuvenation steps. The Augmentation and the Rejuvenation steps are then repeated for each new batch of data.

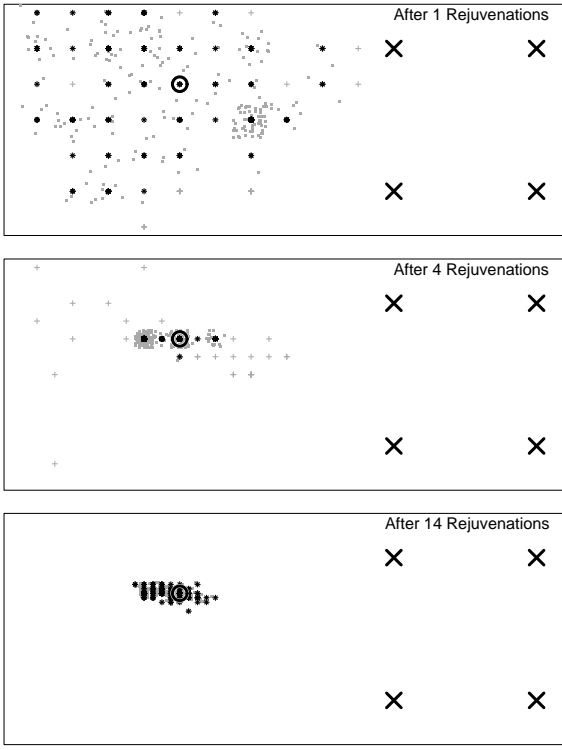


Fig. 1. The annealing and the spatial refinement of the posterior sample of the source location; an open black circle shows the true location of the source, while black stars (and gray crosses) show the spatial-grid in use at each time, and gray stars show 'jittered' posterior location-realizations.

6. A SMALL EXAMPLE

We now demonstrate the impact of annealing and spatial refinement using a small example based on a synthetic setup and data.

The setup consist of a single explosive-type source that is upwind from a array of four sensors; see Figure 1. Synthetic sensor data was generated using a simple Gaussian puff model [1]. Given the first batch of sensor data above the detection limit, we carry out an initial important sampling with 250 particles, followed by MCMC rejuvenation steps that refine the spatial resolution of the source location at 'cooler' temperatures. The evolution of the source location sample is shown in Figure 1 at three iteration points, with only the last one (iteration 14) being conducted at temperature $T = 1$. We note the coarse spatial resolution at high temperatures and the subsequent spatial refinement with cooler temperatures.

Not shown is the augmentation and rejuvenation of the current sample given a new batch of data. In this case, the spatial resolution is kept fixed at the current resolution, but the temperature associated with the likelihood of the new data is selected adaptively, as outlined previously, to yield a good effective sample size.

7. ACKNOWLEDGMENT

This work was funded by LDRD project number 04-ERD-037 and was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

8. REFERENCES

- [1] W.B. Petersen and L.G. Lavdas, "Inpuff 2.0: A multiple source gaussian puff dispersion algorithm — user's guide," Tech. Rep. EPA/600/8-86/024, EPA, 1986.
- [2] J. S. Nasstrom, G. Sugiyama, J. M. Leone, and D. L. Ermak, "A real-time atmospheric dispersion modeling system," in *The Proceedings of the 11th Joint Conference on the Applications of Air Pollution Meteorology*, Long Beach, CA, 2000, The American Meteorological Society, pp. 285–289.
- [3] R.M. Neal, "Annealed importance sampling," *Statistics and Computing*, vol. 11, pp. 125–139, 2001.
- [4] S.N. MacEachern, M. Clyde, and J.S. Liu, "Sequential importance sampling for nonparametric Bayes models: The next generation," *Canadian Journal of Statistics*, vol. 27, pp. 251–267, 1999.
- [5] M.K. Pitt and N. Shephard, "Filtering via simulation: Auxiliary particle filters," *Journal of the American Statistical Association*, vol. 23, pp. 356–359, 1999.
- [6] C. Berzuini and W. Gilks, "RESAMPLE-MOVE filtering with cross-model jumps," in *Sequential Monte Carlo Methods in Practice*, A. Doucet, N. de Freitas, and N. Gordon, Eds., pp. 117–138. Springer, New York, 2001.
- [7] O. Cappé, A. Guillin, J.M. Marin, and C.P. Robert, "Population Monte Carlo," *Journal of Computational and Graphical Statistics*, vol. 13, pp. 907–929, 2004.
- [8] P. Del Moral, A. Doucet, and A. Jasra, "Sequential Monte Carlo samplers," *Journal of the Royal Statistical Association*, vol. 68, pp. 411–436, 2006.
- [9] Jun S. Liu and Rong Chen, "Sequential monte carlo methods for dynamic systems," *Journal of the American Statistical Association*, vol. 93, pp. 1032–1044, 1998.